Naval Research Laboratory
Washington, DC 20375-5320

NRL/FR/5510--04-10,086

Initial Assessment of Human Performance Using the Gaiter Interaction Technique to Control Locomotion in Fully Immersive Virtual Environments

Linda E. Sibert
James N. Templeman
Robert C. Page

Navy Center for Applied Research in Artificial Intelligence
Information Technology Division

Jeremy T. Barron
Justin A. McCune

ITT Industries
Alexandria, VA

Patricia S. Denbrook

Denbrook Computing Services
Centerville, VA

June 30, 2004

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
30-06-2004

2. REPORT TYPE

3. DATES COVERED *(From - To)*

4. TITLE AND SUBTITLE
Initial Assessment of Human Performance Using the Gaiter Interaction Technique to Control Locomotion in Fully Immersive Virtual Environments

5a. CONTRACT NUMBER
N0001403WX20228

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
62235N

5d. PROJECT NUMBER
005

5e. TASK NUMBER
IT-235

5f. WORK UNIT NUMBER
6303

6. AUTHOR(S)
Linda E. Sibert, James N. Templeman, Robert C. Page, Jeremy T. Barron,*
Justin A. McCune,* and Patricia S. Denbrook‡

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Research Laboratory
4555 Overlook Avenue
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER
NRL/FR/5510--04-10,086

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR / MONITOR'S ACRONYM(S)

11. SPONSOR / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
*ITT Industries
Alexandria, VA 22303

‡Denbrook Computing Services
Centerville, VA 20121

14. ABSTRACT
We conducted a first test of the Gaiter locomotion interaction technique to gather data to refine the interaction technique, document system development, and assess system effectiveness. The test showed that the fundamental gesture of stepping in place to walk through the virtual environment worked well. Problems with the side step and back step could be attributed to balance problems caused by the harness. Accurate stopping was more difficult because of the narrow field-of-view of the head-mounted display (HMD). We want the interaction techniques we develop to allow a person to have close to the same ability to coordinate head, arm, and leg movements as in the real world. We use the Iterative Design Process, which includes periodic testing to direct redesign and reimplementation throughout the development process.

15. SUBJECT TERMS
User performance evaluation; Iterative design process; Virtual environments; Infantry simulator; Simulated locomotion

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
Unclassified	Unclassified	Unclassified	UL	23	Linda E. Sibert
					19b. TELEPHONE NUMBER *(include area code)* (202) 767-0824

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

CONTENTS

OBJECTIVE 1

BACKGROUND 1

- Virtual Technology and Environments Demonstration II 1
- Close Quarters Battle 2
- Technical Approach 2
- Gaiter Locomotion Interaction Technique 3

EVALUATION METHODOLOGY 4

- Test Participants 4
- Equipment 4
- Tasks 6
- Data Collection During Task Performance 9
- Test Procedures 9

RESULTS 10

- Time, Accuracy, and Velocity 10
- Rotations 13
- Cybersickness 13
- Presence Questionnaire and Participant Comments 15

DISCUSSION 17

- Conclusions 17
- Changes to the Gaiter Locomotion Interaction Technique 18

ACKNOWLEDGMENTS 19

REFERENCES 19

INITIAL ASSESSMENT OF HUMAN PERFORMANCE USING THE GAITER INTERACTION TECHNIQUE TO CONTROL LOCOMOTION IN FULLY IMMERSIVE VIRTUAL ENVIRONMENTS

OBJECTIVE

The objective of this research was to study the usability of the Gaiter locomotion interaction technique being developed for an individual combatant virtual environment (VE) simulator for training Marines in close quarters battle (CQB) for military operations on urbanized terrain (MOUT). This study is part of a series of tests that will be conducted to evaluate the various interaction techniques developed for the user interface. For example, another interaction technique that allows a person to shoot a rifle will be the subject of a future test. These test results will be used to refine the interaction techniques, document system development, and assess system effectiveness. Immersive virtual environments allow a person to move and interact in a virtual 3-D world. We want the interaction techniques to allow a person to have close to the same ability to coordinate head, arm, and leg movements as people have in the real world. Skills and actions such as aiming a rifle and opening a door should be supported in a way that demands the same skill, timing, and exposure to threats as those actions normally entail in the real world.

Most current usability techniques and guidelines are applicable to a narrow range of interface types such as 2-D graphical user interfaces (GUIs) and have limited use for evaluating VEs and other nonroutine interfaces according to Bowman et al. (2002). One approach that has been effective and is used here is the Iterative Design Process, which is similar to Formative Evaluation described in Hix and Hartson (1993). They define Formative Evaluation as an evaluation of an interactive design as it is being developed to see if it meets its stated objectives and goals. It is more of an exploratory development process used when the design is less fixed. The Iterative Design Process, which we follow, consists of a series of cycles of design, implementation, testing, and redesign based on lessons learned. The tasks used for testing are representative of the larger target tasks for which the interaction technique is being developed. These testing tasks often focus on the critical capabilities needed for the target task. Both quantitative and qualitative data are collected and analyzed. The benefit, in addition to uncovering problems, is that it documents design decisions and the progress made in system development.

BACKGROUND

Virtual Technology and Environments Demonstration II

The Gaiter locomotion interaction technique work is being developed under the Virtual Technology and Environments (VIRTE) Demonstration II Program that is part of the Office of Naval Research's (ONR's) Capable Manpower FNC (Future Naval Capability) and under Naval Research Laboratory (NRL) base funding. The goal of VIRTE Demonstration II is to create a fully immersive simulator for training small unit tactics, techniques, and procedures (TTPs) for CQB. The simulation system will support real-time interaction of small infantry teams as they move and operate within a simulated urban environment. The simulator will be constructed using VE technology, which offers a unique solution and new opportunities to enable individuals and small units to develop competencies needed to win and survive in combat. A need exists for

Manuscript approved January 23, 2004.

a family of realistic, deployable simulators to train skills and rehearse a variety of mission scenarios that are too dangerous, costly, or impossible to practice. These capabilities will become even more important in the future as naval missions become more complex and demanding of human competencies and as live training areas, budgets, and time for training decrease.

Close Quarters Battle

Close Quarters Battle is the term for systematic building clearing. CQB is a complex operation involving a high level of team coordination and individual personal skill. The Marine Corps Institute's manual, "Military Operations on Urban Terrain" (1997), shows building clearing operations using a four- to six-person search party split into a search team and a cover team. The search team clears each room, hallway, and stairway in the building; and the cover team is responsible for security in the immediate vicinity of the search team. The manual diagrams team movements such as door entry. These movements involve shooting well while moving and discerning targets.

Kelly McCann, a VIRTE subject matter expert (SME) on CQB, teaches that CQB is based on initiative-based tactics. The three principles of initiative-based tactics are (1) cover immediate danger areas, (2) shoot threat targets, and (3) protect each other in the team. In other words, while there is an initial plan based on practiced approaches, actual threats are unknown and must be addressed as they are uncovered. Success comes from team interaction. The team members must think for themselves and actively search for an immediate danger area, claim it, move tactically to it, and cover it until it is cleared. Team members come to each other's aid when they think support is needed. The tactics are dynamic. Team members watch each other and respond to immediate danger areas as they evolve.

In a communication to the VIRTE team (2001), McCann emphasized that the principles of initiative-based tactics require a high level of personal skill to proficiently employ dynamic TTPs. A trained team moving to an objective looks choreographed because the skill level is high. The length, quality, and frequency of training are the key elements in skill development.

Technical Approach

The goal of developing a fully immersive user interface to support CQB training is a tall order. It must allow the user to employ the whole body in as natural a way as possible in a localized virtual simulator. The user's natural reflexes to orient, propel, and control the posture of the body should support the user's interaction with the virtual world. Skills and actions, such as aiming a rifle and opening a door, should be supported in a way that demands the same skill, timing, and exposure to threats as those actions normally entail in the real world.

A VE portrays a scene that changes in response to a person's actions to give that person the impression of dealing directly with a 3-D virtual world. This new medium allows people to experience and act within a computer-generated environment. It relies on sensors to track a person's actions, processing to model and render the VE, and displays to present the stimuli to the user. Sensors, processing, and displays are limited in terms of the coverage, resolution, accuracy, and latency with which they can represent forms and actions. All of these can affect the quality of the VE experience.

Locomotion is an essential part of a person's real-world interaction. A person must be able to move naturally through the VE while, in fact, remaining within the bounded physical space of the tracking system. We need to approximate human performance in the real world and to develop designs that are closest to the natural interactions people use in the real world.

Our general approach to interaction technique design is based on principles derived from an understanding of human perception and motor control. We use sensory substitution and motor substitution that maintain a similarity of action, compatibility of interaction, and effect. We want to elicit a person's natural behavior while executing coordinated movements involved in looking, moving, and shooting while performing goal-directed tasks and responding to surprise events.

Sensory substitution involves presenting stimuli that allow a person to perceive the same basic cues as in the real situation. For example, sensory substitution can be used to convey virtual contact. Virtual contact is defined as the interactive simulation of physical contact to allow people to interact with purely virtual objects in ways similar to how they interact with physical objects in the real world. Virtual contact could be presented visually (highlighting the area contacted), through an auditory cue (a bumping sound), or through haptics (feeling a vibration near the site of contact). Sensory substitution plays a secondary role in the development of the Gaiter locomotion technique.

Motor substitution, developed in our laboratory, is the primary principle behind the Gaiter locomotion interaction technique. With motor substitution, the user's experience in the VE derives from what the user's actions are in the real world. Actions can be portrayed either one-for-one with what the user does in the real world or they can be a gestural action based on a real-world action. The effect of the user's gestural action is to produce an illusion of the real-world action in the VE. Gaiter provides a natural, direct means of allowing a user to move over long distances in a VE, while remaining in the relatively small physical space. Gaiter substitutes stepping in place for walking in the virtual world. The simulation of walking compensates for the physical limitation of the tracked area.

Gaiter Locomotion Interaction Technique

The traditional definition of an interaction technique, which was developed for 2-D desktop interfaces, is a way of using an input device to enter information into a computer (Foley et al. (1990)). Interaction techniques are the primary building blocks that craft the user interface. It is remarkable how well the definition holds for 3-D interfaces such as immersive VEs. For these immersive interfaces, a person's body is the device mechanism by which the person interacts with (and thus provides information to) the graphical virtual world (a computer program).

The Gaiter locomotion interaction technique uses the movement of the user's legs to determine the direction, extent, and timing of the user's movement through the VE. Tying virtual locomotion to leg motion allows a person to step in any direction and control the stride length and cadence of virtual steps. Since Gaiter uses only the legs and pelvis, it does not interfere with actions performed by other parts of the body. Furthermore, Gaiter operates in the appropriate coordinate frame of reference — the direction of knee movement with respect to the surrounding environment. Other parts of the body (head, shoulders, arms, torso, etc.) can be aligned as they are in natural postures and movements. This allows people to make use of their reflexes to direct their body in response to external stimuli. People are free to turn their head to look in any direction and to use their hands to manipulate and point at virtual objects in the virtual environment. It also allows users to naturally intermix a wide range of postural motions (e.g., crouching, jumping, and bending to look around objects) with gestural stepping motions.

The attributes of in-place stepping are used to control the characteristics of natural gait as portrayed by the VE. The height and rate of in-place steps map into the stride length and cadence of the steps taken by the avatar, which is the user's virtual body. The system can be tuned to preserve metrics between physical and virtual space, as well as to match specific attributes of natural locomotion, such as perceived velocity and caloric expenditure. Moreover, the pattern recognition system does not have to wait for a gesture to be

complete before responding to knee movements. Virtual motion in response to both actual and gestural steps occurs as the step is taken.

Tying optic flow directly to leg movement makes the Gaiter interface feel like a simulation of walking rather than an indirect control over locomotion. The effect is to create the interactive illusion of walking since the cadence of the stepping movements is reflected in what the user sees.

EVALUATION METHODOLOGY

Test Participants

The test of the Gaiter locomotion interaction technique was run in the Immersive Simulation Laboratory at NRL in May 2002. The ten subjects who participated were employees or contractors of NRL; nine were male and one was female, all between 22 and 35 years of age. The subjects carried the rifle right handed. Seven required no vision correction and three wore contact lenses. Volunteers who wore glasses were excluded because the current head-mounted display (HMD) does not accommodate glasses comfortably. All of the subjects were heavy computer users (from between 40 to 80 hours per week). One subject reported playing first-person shooter video games often; two playing occasionally, two playing infrequently. The remaining four never played them. Two subjects were expert at more than 16 games, two were expert at one or two, and the rest were not expert at any game. Three subjects had military experience. None had ever trained in a simulator.

In general, the subjects were a good group to carry out the initial assessment of the Gaiter locomotion interaction technique. They were similar in age to the Marines who would train on the simulator, but as computer professionals, they were knowledgeable about technology (although not necessarily VE technology) and could offer technical insights.

Unfortunately, a cable was not plugged in for four subjects, which meant a time stamp signal was lost and their path data were unusable. However, their opinion and simulation sickness data were valid because they completed the entire protocol. Because of nausea, a fifth subject withdrew at the beginning of the last path task entitled “Move Around a Curved Path.” The subject did not complete the questionnaire data but the subject’s earlier data were valid and useful.

Equipment

The components of the simulation system used by the subjects were mainly commercial off-the-shelf and commonly used in VE or for marksmanship training. The VE set-up is shown in Fig. 1. The HMD was a V8 model from Virtual Research Corporation. It had a limited but acceptable 48-deg horizontal by 36-deg vertical field-of-view (FOV) and 640 by 480 pixel resolution with one microdisplay per eye. The rifle was an M4 replica AirSoft model by Tokyo Marui Manufacture. It was not loaded with pellets and had been instrumented by an AirSoft gunsmith to send a signal when the trigger was pulled. The trigger-pull signal was converted to digital using a custom board designed in-house and then transmitted to the Gaiter computer over a low-powered radio frequency wireless transmitter. The flak jacket was the standard issue worn by Marines in combat from Point Blank Body Armor. The passive optical tracking system was a Vicon 512 Motion Capture System with 11 cameras. Retro-reflective markers were attached to the HMD, flak jacket, and Velcro bands that were worn on the subject’s legs. A ring of infrared LEDs surrounded the lens of each camera and strobed at 120 Hz. The optical tracking system provided position and orientation data to the Gaiter application. Gaiter also used data from force sensing resistors placed in the subject’s shoes. The sensors were commercial off-the-shelf from Interlink and have been used by the biomedical community to study human gait. The sensor data signal was converted to digital using a custom board designed in-house

Fig. 1 — PVC pipe harness, Vicon tracking system cameras and markers, and V8 HMD shown with the previous wooden frame

and sent to a computer over the low-powered radio frequency wireless transmitter. The flak jacket worn by the subject was attached to a harness, which was part of a custom centering system designed to suspend the HMD cable above so that the subject would not trip over it and to keep the subject in the center of the tracked area. We have found that an untethered person stepping in place with their eyes closed will drift forward. The flak jacket had a Velcro closure that allowed the subject to be quickly disengaged from the harnessing system.

The harnessing system was comprised of a three-legged frame 8 ft-3 in. high with arms 7 ft long. The area within the frame corresponded to the tracked area, approximately $8 \times 8 \times 8$ ft. The harness itself was an initial design made out of PVC pipe at the waist and above the head, with rope connecting the pipes. A centered PVC pipe joint on the top pipe guided a rope and fastener, which connected to the frame. The harness turned about a center point and allowed a few physical steps out from the center.

The image was rendered by an Intergraph ZX10 PC workstation with dual 866 MHz Pentium III CPUs running the Windows 2000 operating system. The scene was presented bioccularly (the same image was presented to both eyes). We did not use stereo images because we could not guarantee that our rendering hardware could generate 30 stereo pairs per second; we could guarantee 30 single image frames per second. We considered smooth, low-latency imagery more important than stereo for the initial tests. The graphic card was a Wildcat 4210, and the image sent to the HMD had 640 by 480 pixel resolution per eye. The graphics application program was in-house custom software built on top of Multigen-Paradigm's Vega. The graphics computer received body position and orientation information from the Gaiter PC over a serial connection. The Gaiter PC was a Dell Precision 530 workstation with dual Pentium 4 Xeon 2.2 GHz processors running the Windows 2000 operating system. It was connected to the Vicon 512 Data Station over Ethernet.

The tests were run using the version of Gaiter custom software that synchronized in-place gestural and virtual steps based on ground contact. The result was that a virtual step was a half step behind its corresponding gestural step. This approach was taken because of its responsiveness, and the effect of the half-step lag

seemed minimal. The results of these tests will point out where improvement should be made in the design of the locomotion interaction technique.

Tasks

The experimental tasks were designed to evaluate how well people can use the Gaiter locomotion interaction technique to perform the basic walking tasks used everyday as well as in CQB operations. In the real world, it is common to walk down hallways and stop to look through doorways; when outside, people walk on sidewalks and stop at intersections. We developed a series of virtual room models showing different tape configurations on the floor that formed paths and stop/pause points. We also studied rotation in the real and virtual worlds by having subjects turn in place a certain number of degrees in both.

The experiment adopted walking along a path marked out on the ground for three reasons: (1) it is an "ecologically valid" task (i.e., one that people perform every day in the real world, such as walking along a sidewalk); (2) if walls had been used to delineate the paths, then collision detection and handling would have impacted task performance, and since collision detection and handling were going to evolve in future VE implementations, we wanted to factor out its impact and focus on locomotion; and (3) it should be easy to mark out pathways using tape on the floor in the real world so that anyone with a long-range tracking system could carry out a real-world comparison study.

All of the virtual tasks took place in a large virtual room that was 28 by 40 ft. The front wall was painted green as a point of reference, and the remaining three walls were beige. Facing the green wall meant that the subject was facing forward. In these virtual rooms, there were pieces of furniture along the walls such as a bookcase, a door, and pictures. The floor was carpeted and the lighting was uniform. Virtual "tape" marked out "corridors" on the floor. Shown in Fig. 2 are the three layouts for the four path tasks. The rectangular path layout had an outer dimension of 12 by 24 ft; the long sides were 3-ft wide, and the short sides were 2-ft wide. The straight path layout ran the length of the room, but the task only required walking within the 24 ft in the middle; the path width was 3 ft. The curved path layout was 24 by 48 ft, and the path width was 4 ft.

The instructions for all of the path tasks were the same, namely to "walk around the path as fast as you feel comfortable while staying near the center of the path." For all of the tasks, including rotations, the subject held a rifle that was slung around the neck on a tactical sling. The tactical sling allowed the subject to release the rifle to adjust the HMD if needed. The subjects were asked to pull the trigger to mark events — start, when aligned with another corridor, and stop.

We had the subjects do a "Rotations" task in the real world (unimmersed standing in the center of the harnessing system in our lab) and in the VE (immersed wearing an HMD in the center of the room with the rectangular path layout). The number of degrees to turn was called out by the experimenter. The subject was instructed to pull the trigger of the rifle before starting the turn, turn that number of degrees, and pull the trigger again when stopped.

All of the path tasks took place in the virtual world with the subject wearing the HMD. The first path task, "Advance Around the Outer 'U'," used the rectangular path layout. The subject started from the back left corner facing forward (toward the green wall). The subject pulled the trigger to start, walked forward around the outer 'U', and pulled the trigger to stop in the back right corner. The subject would then realign their body and do the same thing in reverse from the back right corner. Each subject traversed the path six times, three in each direction. The purpose of this task was to capture how accurately a person could walk forward using the Gaiter locomotion interaction technique. Figure 3(a) shows the view of the virtual room from the start position.

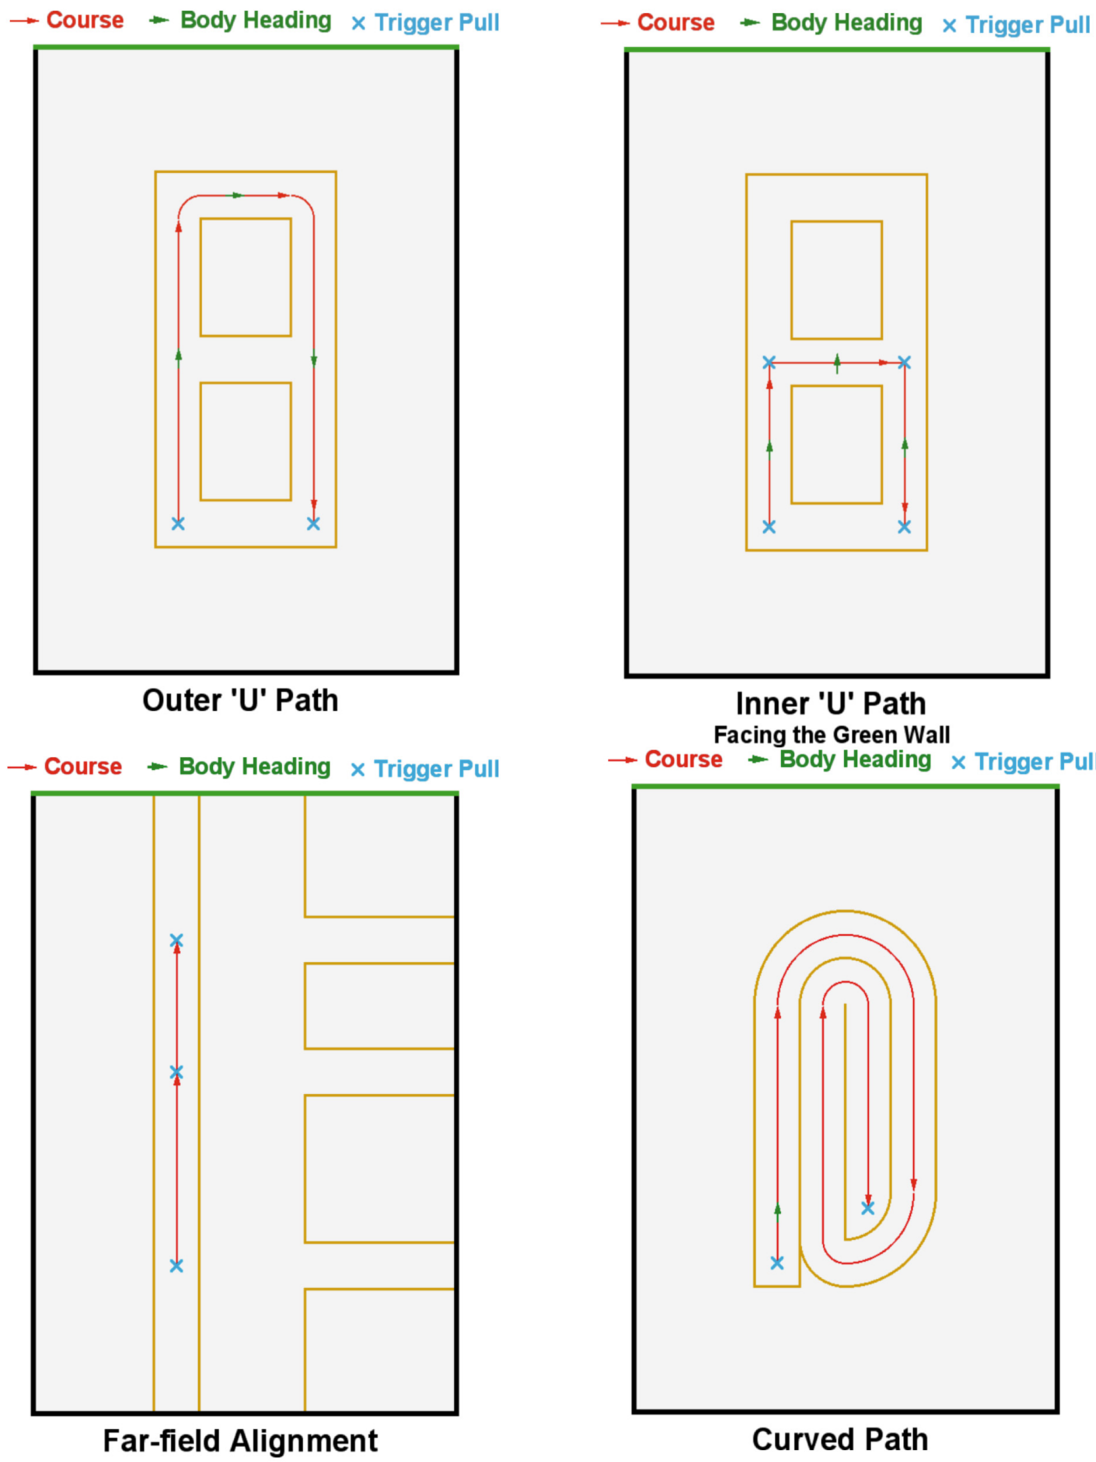

Fig. 2 — Path layouts for the path following tasks

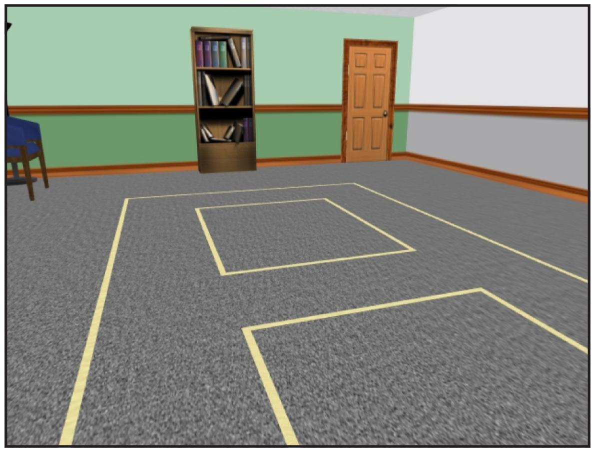

(a) Advance Around the Outer ‘U’ and Move Around the Inner ‘U’

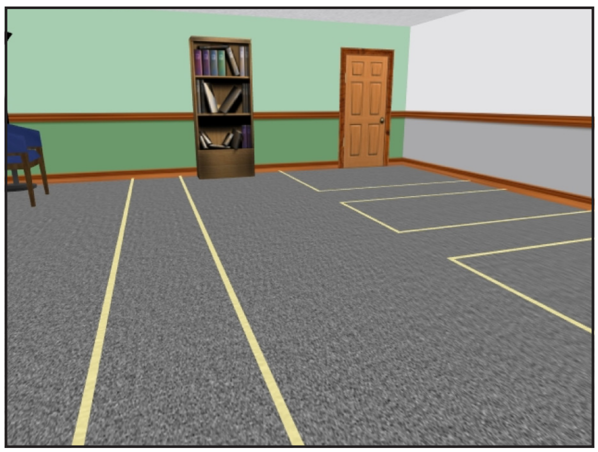

(b) Register Far-Field Alignment

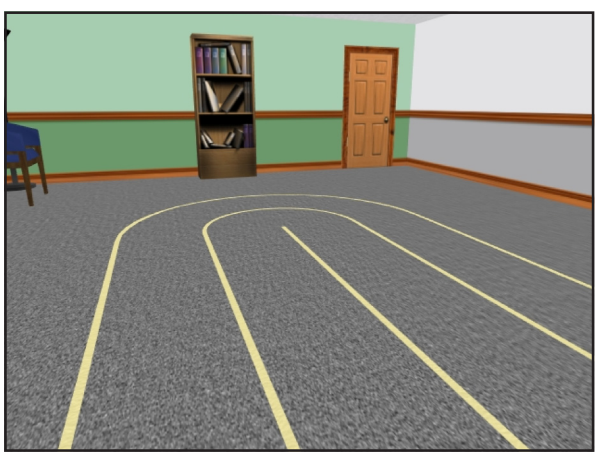

(c) Move Around a Curved Path

Fig. 3 — View of rooms from start position

The second path task, "Move Around the Inner 'U'," used the rectangular path layout. The subject again started from the back left corner facing the green wall and alternated six times back and forth to the back right corner. This time, the subject always faced the body forward. That meant that a subject would walk forward along the first side, side step on the second, and walk backwards on the third. The subject pulled the trigger to start and stop, and stopped and pulled the trigger when in the center of each path intersection. The purpose of this task was to test the accuracy of the different stepping patterns — forward, side step, and back step — and also to test the person's ability to stop with intention at a location. Again, Fig. 3(a) shows the view of the virtual room from the start position.

The third path task, "Register Far-Field Alignment," used the straight path layout. The subject started opposite the center of the intersection with the last corridor furthest from the green wall. The subject pulled the trigger to start (opposite the center of the corridor furthest from the green wall) and stop (opposite the corridor nearest the green wall), and also at the center of the middle corridor. The subject did not have to stop at the center of the middle corridor, only to pull the trigger when the body was aligned with it. The subject could move along the straight path using any of the stepping patterns. The subject performed the task three times in one direction and three times in the other. The purpose of the task was to test the person's ability to stop with intention using a distant cue. Figure 3(b) shows the view of the virtual room from the start position.

The last path task, "Move Around a Curved Path," used the curved path layout. Again, the subject started from the back left corner facing the green wall and pulled the trigger to indicate start and stop. Stop this time was in the center of the spiral. The subject traversed the path six times, three in each direction. The subject was instructed to move forward along the path, which caused the person to turn. The purpose of this task was to test a person's ability to accurately walk forward while turning. Figure 3(c) shows the view of the virtual room from the start position.

Data Collection During Task Performance

Data were collected to support the calculation of basic metrics associated with locomotion, namely accuracy (the ability to stay in a path or to stop) and rate of movement (velocity). The time-stamped position and orientation of the head, center of gravity (COG), and rifle were stored as rapidly as possible using a binary representation. COG was the approximate x and y center of the person calculated from the position of the rigid object defined by the markers on the flak jacket. Data were written to disk at the frame rate, which averaged 30 frames per second. A series of analysis programs were written to calculate the performance measures. For Rotations, the measures were error (in degrees), time from start to stop trigger pulls, and translation from the center using the subject's center of gravity. For the path tasks, they were path accuracy (the root-mean-square deviation from a fixed centerline), terminal accuracy (the distance at an endpoint), and velocity (the distance along a path divided by time).

Test Procedures

A major problem with current HMDs is the limited FOV (in our case, 48-deg horizontal by 36-deg vertical). To improve a subject's ability to view the room and see the path, we compressed a 72-deg horizontal by 54-deg vertical geometric FOV into the 48-deg horizontal by 36-deg physical FOV. From informal testing in our laboratory, we found that the path following tasks required excessive head motion to see the scene when the image was displayed uncompressed (i.e., one-for-one). Psotka et al. (1998) demonstrated that a small frame limiting a person's FOV of an accurately projected image causes distortion effects on self-location, which supports our informal findings. A small frame places the viewer closer to the object of view as occurs with movies. Even with the compressed FOV, however, it was hard for subjects to adequately see their feet on the path. The 72-deg horizontal by 54-deg vertical FOV provided by the compression was still

quite limited compared with natural vision, which provides about 120 deg per eye horizontal and vertical with 90 deg of overlap. The compression effect is shown in Fig. 4.

We tracked a subset of the body and displayed a self-avatar that had limited movement (again, an avatar is the image of the person in the virtual environment). Tracker markers were placed on the HMD to track the subject's head, the flak jacket for the torso, and the two lower legs. The lower leg data were combined with the data from the pressure sensors to calculate the locomotion displacement. The avatar was full-body but did not have full movement. The head reflected in the viewpoint moved as the person looked around and walked through the VE. Legs and feet were drawn locked together and were shown hanging straight down; the position of the feet indicated the location of the avatar along the path. (An articulated full-body avatar was not available for this experiment.) The rifle was also tracked and could be seen in the VE if it were picked up or extended, but the avatar's arms were not displayed. Because the subjects were instructed to carry the rifle in a low-carry position, it rarely if ever, came into view. The limited appearance and movement of the avatar did not affect the task. The subject could freely move the head and see the feet to gauge position along the path.

The series of tests took about an hour and a half. Before testing, the subjects completed a background questionnaire developed in-house and the Simulation Sickness Questionnaire (SSQ) by Kennedy et al. (1993). The order of tasks was Real World Rotations, Virtual Reality Practice, VE Rotations, Advance Around the Outer 'U', Move Around the Inner 'U', Register Far-Field Alignment, and Move Around a Curved Path. The subjects rested between tasks but were taken out of the VE for at least 5 min after the VE Rotations and the Move Around the Inner 'U' task. Once a subject had completed the tasks and had taken off the equipment, the subject completed a second SSQ, the Presence Questionnaire described in Witmer and Singer (1998), and an oral debriefing. If the subjects felt any ill effects, they were asked to stay until they improved. The subjects were offered water throughout the tests; and cookies and crackers were available.

RESULTS

Time, Accuracy, and Velocity

Path accuracy (the root-mean-square (RMS) deviation from a fixed centerline) was calculated for the last two thirds of the path for both straight paths and curved segments. The idea was, for straight path tasks, to eliminate the turning after the cornering acceleration phase and to capture the subject moving at stride. In the curved path, the point was to isolate the subject's ability to maneuver around curves. Terminal accuracy (the distance at an endpoint) was Euclidean distance from the true center of the intersection of two paths or the extension of a far-field path with a path. Velocity (the distance along a path divided by time) was calculated for the same areas of the path used to calculate path accuracy. The values were calculated for each subject for each task and put in a spreadsheet. The purpose of this part of the evaluation was to look for patterns in the data that point to problems using the Gaiter locomotion interaction technique. During testing, the experimenter kept notes of any problems with the system, for example, if there was a jump in the display or if the subject was struggling with the harness. These events were highlighted in the data. Table 1 is a summary of the mean and standard deviation values.

Figure 5 shows a representative trial for each type of path task, displaying position data from the COG for one subject on one traversal. The subject's awkwardness with backstepping is evident in the Terminal Accuracy diagram (the right path) as is the difficulty turning along the curved path without cutting corners.

The most basic observation was that there was a large degree of variation in the way first-time users controlled the Gaiter locomotion technique — some were very careful while others were cavalier. The subjects who played a lot of computer games seemed to play the system and seek some performance goal. In

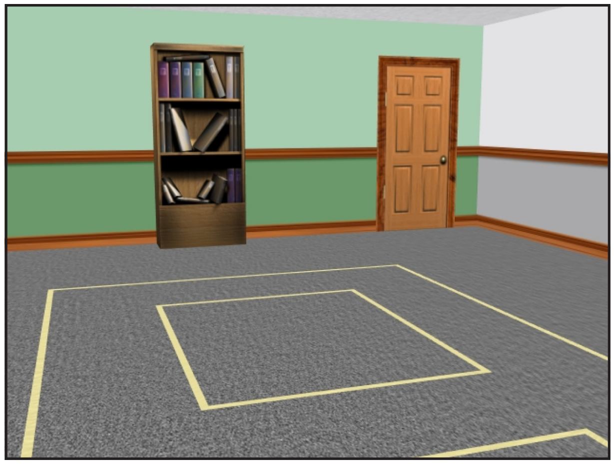

One-to-One FOV
48-deg horizontal by 36-deg vertical physical FOV
48-deg horizontall by 36-deg vertical geometric FOV

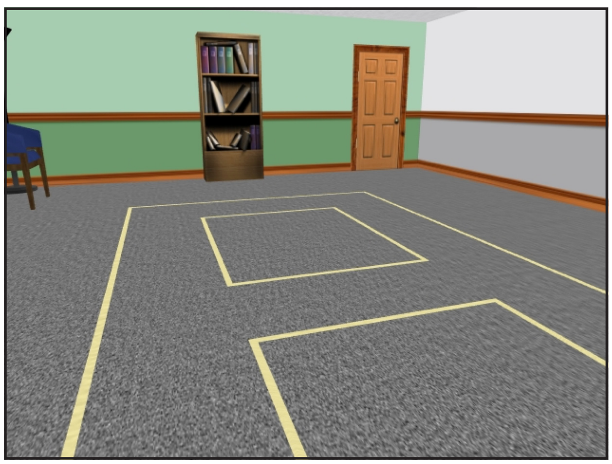

Compressed FOV
48-deg horizontal by 36-deg vertical physical FOV
72-deg horizontal by 54-deg vertical geometric FOV

Fig. 4 — Effect of compressed FOV on image

Table 1 — Summary Results (Forward Values Highlighted in Red)

		Path Accuracy (RMS in ft)			Velocity (ft/s)			Terminal Accuracy (ft)	
		Mean	*Std.Dev.*		Mean	*Std.Dev.*		Mean	*Std.Dev.*
Outer 'U'	Forward	0.50	*0.42*	F	5.09	*2.67*			
Far-field	Forward	0.52	*0.46*	F	4.93	*1.63*	F	1.09	*0.76*
Inner 'U'	Forward	0.35	*0.21*	F	4.56	*1.31*	F	0.99	*0.64*
	Sideways	1.30	*0.89*	S	1.62	*0.67*	S	1.67	*1.12*
	Backward	0.53	*0.37*	B	2.73	*0.93*	B	1.34	*0.89*
Curve	Advance	0.97	*0.35*	A	3.22	*1.24*			

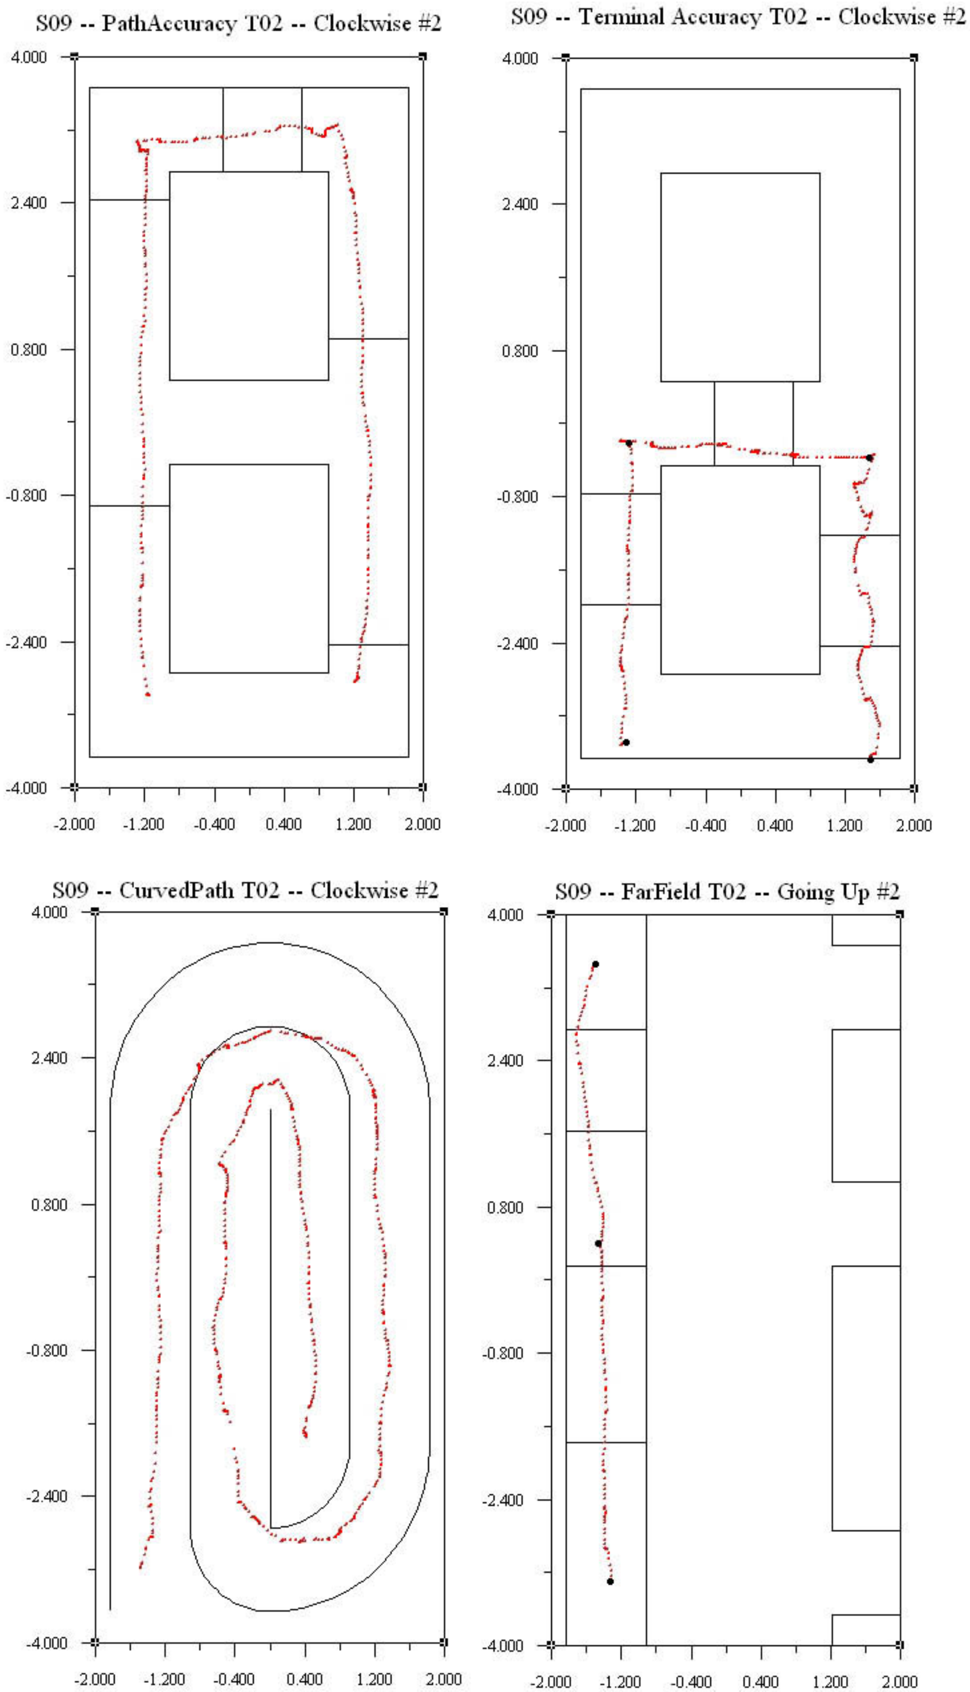

Fig. 5 — Example path trajectories (dots indicate trigger pulls)

general, from data analysis, there was a consistent accuracy and speed associated with walking forward down a 3-ft path regardless of the task. Subjects walked forward at an average velocity of about 5 ft/s while achieving a path accuracy of about 0.5 ft (RMS) and a terminal accuracy of about 1 ft. Some subjects indicated that they had problems balancing with the side step because of the harness. Some indicated having trouble seeing their feet because of the narrow FOV of the HMD, which can explain the variation in the terminal accuracy results. The average velocity for side step and back step was faster than for forward steps. One explanation could be that the formula to calculate distance traveled for each step moves a person too far. It is important to tune the system to match real-world performance.

Rotations

The subjects performed six rotations, first in the real world, and then in the VE. The subject was tracked in both conditions. For real-world rotations, the subject was suited in all of the equipment needed for tracking, except that a headband with markers was substituted for the HMD. The subject was not attached to the harness. For virtual rotations, the subject changed to the HMD and was attached to the harness. This task was simple but served two purposes. First, it introduced the equipment and allowed the subject to practice moving while wearing the harness. The more important reason was that it provided a measure of how wearing a HMD and turning with the inertia of a harness affected performance (ability to stop accurately, time to complete rotation, and translation error).

Unfortunately, the set of rotation angles was not held constant across conditions so formal analysis was not possible. An exploratory analysis suggested that the HMD and harness affected rotation. The results from the two 180 deg rotations were used as the comparison. It took longer to turn in the VE (an average of 5.24 s in the real world vs 6.60 s in the VE), there was more translation (1.12 ft in the real world vs 1.47 ft in the VE), but the error in degrees was less in the VE (3.72 absolute deg in the real world vs 2.14 deg).

The dominant factor seemed to have been the harness. There was inertia to overcome to begin a turn, the harness would swing slightly as the subject turned, and swing more as the subject stopped. The fewer degrees of error might have occurred because there were stronger direction cues in the VE than in the real world. The furniture and decorations were less cluttered and more distinct in the VE while in the real world, the frame was off-center to the walls, and the room was filled with random cables and equipment. The results are suggestive, and more research is needed to determine if any of these differences are significant.

Cybersickness

For some people, being immersed in a VE can induce discomfort because of cybersickness, which has effects similar to motion sickness such as fatigue, dizziness, awareness of the stomach, and increased salivation. VE technology is an approximation of the real world. A person in a VE system views the virtual world through an HMD, which has a limited FOV (in our case, 48-deg horizontal by 36-deg vertical). In order to see the virtual world, a person must crane the neck unnaturally to see straight down, up, left, or right, which can lead to stress and unease. The resolution is limited (640 by 480 for the V8 used for these tests) and the image does not appear sharp, which can induce eyestrain. The system can also have lag so that when a person moves the head, the correct image appears a fraction of a second later, throwing off hand/eye coordination.

We used the SSQ developed by Kennedy et al. (1993) for military flight simulators to measure sickness. It is a commonly used measure for simulator sickness. There are 16 questions that ask about the severity of symptoms (none, slight, moderate, or severe). The maximum possible *raw score* is 48. The answers are also weighted to obtain three *subscales* [nausea (N), oculomotor problems (O), and disorientation (D)] and a total score. Symptoms making up the nausea scale are general discomfort, increased salivation, sweating, nausea, difficulty concentrating, stomach awareness, and burping. Symptoms making up the oculomotor

problems scale are general discomfort, fatigue, headache, eyestrain, difficulty focusing, difficulty concentrating, and blurred vision. Symptoms making up the disorientation scale are difficulty focusing, nausea, fullness of head, blurred vision, dizzy with eyes open, dizzy with eyes closed, and vertigo. The SSQ was given before and after the subject participated in the VE tests. The results are from the nine subjects who completed all the tests.

The calculated SSQ *total score* was 20.8 (total score is a calculated *weighted* score different from the raw score). To put that figure in perspective, the average SSQ total score for the eight VE experiments using four different systems reported by Stanney et al. (1997) was 29. The individual scores for those eight experiments had a broad range of 16 to 55. In comparison, our total score was fairly low. The factor that raised our total score the most was "sweating." The *raw score* for sweating accounted for 10 points out of a total of 38, and it was the highest-scored symptom of the 16, with the other 15 symptoms averaging just 1.9 points each. Sweating was inevitable since the subject wore a flak jacket and carried a rifle while doing the physical act of walking. The protocol was to keep the room cool, but sweating was expected. No subject reported sweating before the tests. After the tests, only two subjects reported no sweating, while four reported slight sweating and three moderate sweating. None reported severe sweating. Much of the sweating was likely due to the subjects' body temperature rather than the effect of cybersickness.

The effect of sweating was apparent in the calculated subscale scores, which were nausea (27.56), disorientation (15.47), and oculomotor problems (11.79), giving the pattern N>D>O. Later work by Stanney, et al. (1997) looked at the SSQ results from their study of the eight VE experiments using four different systems and compared them to a database of reports from flight simulators. They concluded that the VE systems presented a different pattern of results than those for flight simulators. Stanney and Kennedy (1997) reported that simulator sickness scores for flight simulators have the pattern O>D>N, while VE systems have D>N>O. Airsickness and seasickness have the pattern N>D>O. In addition, the amount of overall sickness was higher for VE systems than for flight simulators. One reason suggested to explain the lower incidence of sickness with flight simulators was that people who become pilots are less inclined to suffer from motion sickness (self selection) and pilots who train in military flight simulators are often motivated not to report medical symptoms from their use. Stanney et al. (1997) called VE sickness "cybersickness" to distinguish it from the simulator sickness associated with flight simulators. The pattern of subscales obtained in our study of virtual locomotion of N>D>O is the same as the pattern reported for airsickness and seasickness, not the VE pattern of D>N>O. The difference is in the artificially high value of the nausea scale, which changes the results from the expected VE pattern.

Stanney and Hash (1998) showed that either an *active control* to move through the VE or an *active-passive control* mapped to the needs of the tasks were effective methods for minimizing cybersickness. Active controls allow users to control their movements within a VE. Active-passive controls further constrain the control in a task-specific manner. The Gaiter locomotion interaction technique allows a person to control their movement through the VE at a rate and with a similar ability to coordinate looking, moving, and manipulation as afforded in the real world. The subscale scores for our test (nausea 27.56, disorientation 15.47, and oculomotor problems 11.79) were more like the active-passive control in the Stanney and Hash (1998) paper, again remembering that our nausea score is higher because it contains the sweating factor. The active-passive scores from Stanney and Hash (1998) were nausea (11.53), disorientation (13.12), and oculomotor problems (16.24); the scores they reported for active controls were nausea (19.08), disorientation (26.85), and oculomotor problems (31.32). Gaiter seems to behave as an active-passive control in that the technique is constrained for performing real-world tasks.

The background questionnaire included questions about the person's previous susceptibility to motion sickness and current well-being, in addition to demographic information reported in the section on Test Participants. The questions were included to see whether they could be used to screen for people susceptible

to cybersickness. The results suggested that previous car or motion sickness is a somewhat useful question for screening participants. The only two subjects who reported previous difficulty with car or motion sickness had higher SSQ total raw scores (two and three before the VE and five and nine after, out of a maximum of 48 points) than the overall average (1.3 before and 4.1 after). But the score of nine, the highest score after the tests, was also scored by a subject who indicated no previous problems with car or motion sickness.

It was found that current tiredness and previous sickness with first-person shooter games were only moderately useful for screening. There were no strong patterns. Three subjects indicated tiredness. Of the two subjects who reported moderate tiredness (a score of two), one had only an average score on the SSQ after VE (a one) and the other scored second to the highest (a six). Another person who reported only slight tiredness scored at the average on the SSQ given after testing (a four). Only one subject reported any problems with sickness while playing first-person shooter games (at the moderate level). That subject's SSQ scores were also slightly above the average before and after testing (a two and five).

But even the highest raw score of nine (out of 48) after VE indicated only minimal symptoms. The highest score before VE was a three (again out of 48), which is rated as negligible symptoms. There is no indication that any of these subjects should not have participated.

The subject who withdrew from the "Move Around a Curved Path" task (the last task) because of general discomfort and nausea did not report problems with car or motion sickness or while playing first-person shooter games. (This subject was not included in the SSQ analysis above.) However, the subject's before-SSQ score was high (14). The experimenter suggested that the subject should not participate, but the subject expressed a strong desire to participate and was not worried because of being a frequent first-person shooter game player. Later, the subject mentioned becoming extremely seasick repeatedly in the past. In hindsight, the subject should have not been allowed to participate even though the level of sickness when he withdrew was not overly severe and improved by resting.

In general, the results show that the VE simulator using the Gaiter locomotion technique did not produce high sickness values even with sweating included. If sweating were discounted, the total and raw scores would be even lower. The VE simulation system still suffered from an HMD with a small FOV, which interfered with the task. Part of the task was to stop at the intersection of two paths, and the subjects wanted to look down at their feet. Because of the small vertical FOV, many subjects were seen to crane their necks down unnaturally, which could cause discomfort. The harness was also seen to interfere with a subject's ability to make stepping gestures, which could result in stress. We anticipate the possibility of a lower SSQ total score with improved technology.

Presence Questionnaire and Participant Comments

We used the Presence Questionnaire by Witmer and Singer (1998) to gather information about the subject's experience in the VE. As defined in Witmer and Singer (1998), *presence* is "the subjective experience of being in one place or environment, even when one is physically situated in another." The usefulness of the Presence Questionnaire for us was that it measured not only the extent of a subject's involvement in the VE, but also the effect of the various components of the VE (such as the Gaiter locomotion interaction technique) within that experience. It asks questions about a subject's subjective experience in the VE and how well the subject was able to control actions. The current simulation system did not include audio, objects moving through space (except for the rare glimpse of the rifle), touch, or walking around a stationary object. Questions on those topics were not counted.

Three scale scores (*Involved/Control*, *Natural*, and *Interface Quality*) and a *Total* score were calculated. *Involved/Control* included questions about how well a subject could control events and how involving the

subject found the visual aspects. *Natural* included questions about the naturalness of the interaction techniques including how well the subject could control movement in the VE. *Interface Quality* included questions about whether the interaction techniques or the display itself distracted the subjects from performing the tasks and the extent the subject felt able to concentrate. The questionnaire uses a seven-point scale scored one to seven and is anchored at each end and at the midpoint.

The *Total* score for this test (based on 16 questions) was six. The *Involved/Control* (three questions) and the *Interface Quality* (14 questions) scores were also six each. The *Natural* score (three questions) was four. Six is near the top of the positive end of the scale, and it indicates that we are doing very well. Four is at the midpoint, which, although not negative, indicates some need for improvement.

The subject's comments during an informal debriefing at the end of the session put these scores into perspective. The subjects were an excellent group to perform the first tests of the locomotion interaction technique because they have had extensive experience with computers, many have played video games, and three had military experience.

A major complaint was about the harness. The subjects could feel it pull at them, and several reported that it was hard to know where the center was located (there is slack over a 1.5 ft diameter of movement). The harness was the reason most subjects gave for not feeling stable while doing the side step. It would also apply torque to the subject when the subject was off center during the "Move Around the Curved Path" task. Some subjects also had trouble with the back step and felt it was not natural, in part because of balance problems. However, a few subjects mentioned that back stepping in the real world is the least natural step. Most subjects thought that stepping in place to walk forward worked well and was intuitive.

The subjects complained that the HMD was front heavy, and one subject particularly noticed the force and weight of the HMD while turning. Most subjects reported problems gauging where they would stop. Part of the problem, they said, was the poor vertical FOV and the difficulty looking down. But they also reported that problems with stopping resulted from not being able to control and predict the length of a step.

Subjects mentioned that they found it odd that the feet and legs just hung straight down and did not move. They felt the model of the room was good. A few subjects noticed some lag in system response, particularly those who had brought the rifle into view.

One problem mentioned by many subjects was hearing directions from the experimenter outside of the VE. Several subjects found this disconcerting and mentioned that it would be better to display an avatar or some other image representing the person in the virtual room rather than having a disconnected voice.

We asked the subjects about their impression of the FOV in the HMD and most thought the horizontal FOV was not too different from the 180-deg FOV in the real world. Moreover, they felt the tasks did not require a wide horizontal FOV. Most felt, however, that the vertical FOV was far less than the 120 deg available in the real world. The impression was that the subjects either did not notice or feel they had any adverse effects from our compressing the FOV to display a 72-deg horizontal by 54-deg vertical image.

In general, the subjects were more conscious of the problems they had with the locomotion stepping patterns than the FOV. However, the poor vertical FOV seemed to have directly contributed to imprecise stopping. Also, balance and the ease of learning the leg gestures seemed affected by problems with the harness, and the balance and weight of the HMD.

DISCUSSION

Conclusions

Most subjects reported that walking forward by stepping in place worked well. There was a consistent accuracy and speed associated with walking forward down a 3-ft path over the several different paths. The Gaiter locomotion interaction technique successfully solved the problem of allowing a user to move through a VE in a controlled manner similar to walking in the real world while remaining in the restricted tracker area.

Unfortunately, the VE equipment can add artifacts that make it impossible to isolate the effect of the locomotion technique. The biggest problem discovered in the first test of the Gaiter locomotion interaction technique was that the rope and pole harness interfered with the subject's stepping. On the positive side, the harness did serve the purpose of centering the person. A natural tendency when people have their eyes closed and stepping-in-place is to drift forward. The version of the harness used in the experiment, however, interfered with stepping more than expected. It unbalanced subjects while side stepping, and it had a tendency to torque the subject when turning off center, as during the "Move Around a Curved Path" task. The back step was also difficult for some subjects because of difficulty balancing with the harness. Back stepping over a long distance is not something people do often in the real world, and a number of subjects mentioned that as the reason they did not like the back step gesture rather than being unbalanced.

The HMD was a problem. Its nose-heavy design bothered some subjects. The narrow FOV (even using a compressed FOV to display more of the scene in the HMD) made looking down at the feet difficult. The subjects had to crane their necks, and many subjects simply stopped trying, which meant they were less sure of where they stopped when asked to stop in a corner. The FOV of the HMD used in the experiment was current state-of-the-art.

Several subjects mentioned being bothered by the fact that the avatar's legs hung straight down and did not move. The amount of articulation of the avatar is determined by the number of tracked body segments and by how that information is used. For this test, we tracked three segments: the head (by placing tracker markers on the HMD), the torso (from markers on the flak jacket), and the two lower legs (from markers attached to Velcro bands worn by the subject on the legs). We only displayed changes to the subject's viewpoint generated from head movement, real displacement from real steps, and virtual displacement calculated from stepping gestures.

Some subjects noticed lag in the system, particularly when they brought the rifle into view. However, they did not feel the simple tasks in this test were affected by system lag.

Cybersickness for our simulation system was low compared to the VE systems tested by Stanney et al. (1997), which we had expected from the research by Stanney and Hash (1998) that showed that active-passive controls are associated with lower cybersickness scores. The Gaiter locomotion interaction technique is an active-passive control in that it allows a person to control their movement in the VE at a real-world rate and with a similar ability to coordinate looking, moving, and manipulation. The cybersickness score for our system will always be elevated because a user wears a flak jacket and carries a rifle, and the target task of clearing a building is hard work, all of which induce sweating. We will work with the other factors and look for ways to lower the score. One way is to improve the technology, such as using an HMD with a wider FOV and better resolution, minimizing system lag that disrupts the visual response to a user action, and improving the fit of the equipment.

One observation about working with non-military subjects who are not accustomed to wearing a flak jacket and carrying the weight of a rifle is that it is important to look closely at the person's initial SSQ. It is imperative to make it clear to any person who is already a little tired and warm that the VE will require physical work and to impress upon the person that it is not just like a first-person shooter game. There was a sense with a couple subjects that they were surprised about the nature of the system even though it had been explained beforehand. Military personnel who have tried the system during a demonstration have had no problem with the level of physical activity.

Lastly, many subjects mentioned that it was disconcerting to have the experimenter speak to them without a representation of that person in the VE. It broke their sense of immersion and presence. For different tests where immersion and presence are critical, it would be more important to minimize sounds from outside the VE.

Changes to the Gaiter Locomotion Interaction Technique

The first test of the usability of the Gaiter locomotion interaction technique provided useful insights that have led to the redesign and reimplementation of parts of the simulation system. At some point, these changes will be tested in keeping with the Iterative Design Process.

The rope and PVC pipe harness caused more problems than expected. Since the experiment, we have installed a next-generation steel harness developed by Dr. Roger Kaufman of The George Washington University. It has a springy steel mechanism that takes up slack evenly so that a user does not feel any odd jerks or torques when moving off center. The turning elements of the harness have low angular inertia. The harness is easy to put on and take off and uses an adjustable belt, which comes as part of a Marine's standard issue MOLLE gear (a backpack system). The harness can be adjusted for people of different heights and provides more vertical play than previous designs. Informal demonstrations have shown that the new harness does not cause the balance problems previously experienced with the side step and back step.

To solve the HMD problem, we have purchased a new HMD from NVIS, Inc., that has a better ergonomic design, higher resolution, and brighter image. The FOV is the same as the current V8. Even with a restricted FOV, we expect that users wearing it will have a better situation awareness in the VE. We are currently working with two companies on a Phase 2 STTR for wider FOV HMDs, on the order of 110-deg horizontal by 68-deg vertical. These HMDs will have good resolution and brightness using the new microdisplays that are now available, and both companies are emphasizing good ergonomic design. We are hopeful that the next generation of HMDs will improve the quality of a user's interaction in the VE.

It was hard to separate the side effects caused by the harness and the FOV from problems with the stepping gestures. However, some of the problems the subjects had with knowing where a step would stop had to do with the way the Gaiter locomotion interaction technique was implemented. We are investigating the development of an alternate step phasing for the Gaiter. The original control used in the tests synchronized in-place gestural and virtual steps based on ground contact, resulting in a virtual step that was a half step behind its corresponding gestural step. The prospective approach synchronizes the gestural and virtual steps based on leg separation, eliminating the half-step lag and providing more immediate update of the virtual displacement. Discriminating between real and virtual steps is more difficult with the new phasing, however, complicating the recognition of inherently physical actions such as turning in place. We will test to determine the effectiveness of the new approach.

We are also improving the implementation of the Gaiter locomotion interaction technique. We have designed and implemented an alternate means of detecting when the feet make or break contact with the ground by using only the tracked motion of the feet and legs. This allows us to run the Gaiter locomotion

interaction technique without the insole pressure sensors, which complicated the process of suiting up and were prone to breakage. It also eliminates the need to synchronize the pressure sensor data with the data from the Vicon tracker.

We have implemented and tested a full-body (16-segment) avatar to replace the original three-segment avatar. We have experimented with different marker layouts for the arms, hands, and feet, which were not required before, and also with the hip belt attached to the centering frame. A good maker layout is one that is easily distinguished by the tracking system (all marker patterns are distinct from one another and can easily be reacquired when markers are obscured). A full-body marker layout is necessary for a more articulated avatar. We have also developed an initial definition for the visual model of the new avatar and worked with the artist doing the modeling. The visual model needs to have a more accurately jointed and articulated skeleton to move realistically. A better avatar would have helped the subjects' sense of being part of the VE test environment. A functioning avatar provides a sense of personal scale, relative distance, and alignment to virtual objects in the near field, and gives the person a sense of being included in the virtual landscape. Having legs that did not move lessened the reality of the VE.

We have worked to improve latency throughout the system. We have made changes to the code base to improve efficiency and maintainability, in addition to adding new functionality such as the ability to support a full-body avatar. We have replaced the serial interface between the PC processing the Gaiter code and the rendering PC with an Ethernet/UDP protocol, reducing the overall system latency by 6 ms. We have also replaced the Gaiter and rendering PCs and have achieved roughly double the frame rate. Lastly, we are upgrading to the Vicon 612 tracking system, which has cameras with a wider FOV and higher resolution that should give better performance.

The improvements to the harness and system latency are expected to have a positive effect on cybersickness. Future subjects will not have to fight the harness to complete the tasks, and the system will be more responsive. If the rifle is used, a subject will be able to move it and see less visual lag. A better HMD will also increase a subject's comfort so that the subject can forget about the technology and fully concentrate on the tasks. There is little we can do about sweating. A flak jacket is necessary for a realistic CQB simulation, and the task of building clearing takes physical energy. An important lesson from the first locomotion test is to closely assess subjects not familiar to this type of physical exertion and to remind them that the VE simulation is not just another computer game.

ACKNOWLEDGMENTS

This work was sponsored by the ONR 6.2 grant "Mobile Interaction for Dismounted Combat Simulators" and the NRL 6.2 grant "Avatars for Fully Immersive Simulation." The work is part of the VIRTE program, which is part of the Capable Manpower FNC program at ONR. We wish to thank LCDR Dylan Schmorrow, the Program Manager of VIRTE, for his advice and support.

REFERENCES

Bowman, D.A., J.L. Gabbard, and D. Hix, 2002. "A Survey of Usability Evaluation in Virtual Environments: Classification and Comparison of Methods," *Presence* **11**(4), 404-424.

Foley, J.D., A. van Dam, S.K. Feiner, and J.D. Hughes, *Computer Graphics: Principles and Practice* (Addison-Wesley, Reading, MA, 1990), Ch. 8.

Hix, D. and H.R. Hartson, *Developing User Interfaces* (John Wiley & Sons, Inc., New York, 1993), Ch. 10.

Kennedy, R.S., N.E. Lane, K.S. Berbaum, and M.G. Lilienthal, 1993. "Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness," *The International Journal of Aviation Psychology* **3**(3), 203-220.

"Military Operations on Urban Terrain," Marine Corps Institute Report MCI 03.66b, August 1997.

McCann, K., Crucible, Fredericksburg, VA, private communication as VIRTE SME, July 2001.

Psotka, J., S. Lewis, and D. King, 1998. "Effects of Field of View on Judgements of Self-Location: Distortions in Distance Estimations Even When the Image Geometry Exactly Fits the Field of View," *Presence* **7**(4), 352-369.

Stanney, K.M. and P. Hash, 1998. "Locus of User-Initiated Control in Virtual Environments: Influences on Cybersickness," *Presence*, **7**(5), 447-459.

Stanney, K.M. and R.S. Kennedy, 1997. "The Psychometrics of Cybersickness," *Communications of the ACM*, **40**(8), 67-68.

Stanney, K.M., R.S. Kennedy, and J.M. Drexler, 1997. "Cybersickness is not Simulator Sickness," *Proceedings of the Human Factors and Ergonomics Society*, 1138-1142.

Witmer, B.G and M.J. Singer, 1998. "Measuring Presence in Virtual Environments: A Presence Questionnaire," *Presence* **7**(3), 225-240.